arXiv:cs/0408049v1 [cs.NE] 21 Aug 2004

Using Stochastic Encoders to Discover Structure in Data*

Stephen Luttrell

February 1, 2008

Abstract: In this paper a stochastic generalisation of the standard Linde-Buzo-Gray (LBG) approach to vector quantiser (VQ) design is presented, in which the encoder is implemented as the sampling of a vector of code indices from a probability distribution derived from the input vector, and the decoder is implemented as a superposition of reconstruction vectors. This stochastic VQ (SVQ) is optimised using a minimum mean Euclidean reconstruction distortion criterion, as in the LBG case. Numerical simulations are used to demonstrate how this leads to self-organisation of the SVQ, where different stochastically sampled code indices become associated with different input subspaces.

1 Introduction

In vector quantisation a code book is used to encode each input vector as a corresponding code index, which is then decoded (again, using the codebook) to produce an approximate reconstruction of the original input vector [1, 2]. The purpose of this paper is to generalise the standard approach to vector quantiser (VQ) design [3], so that each input vector is encoded as a vector of code indices that are stochastically sampled from a probability distribution that depends on the input vector, rather than as a single code index that is the deterministic outcome of finding which entry in a code book is closest to the input vector. This will be called a stochastic VQ (SVQ), and it includes the standard VQ as a special case. Note that this approach is different from the various stochastic approches that are used to train VQs (see e.g. [4, 5, 6]), because here the codebook itself is stochastic, so the use of probability distributions is essential both during and after training.

One advantage of using the stochastic approach, which will be demonstrated in this paper, is that it automates the process of splitting high-dimensional input vectors into low-dimensional blocks before encoding them, because minimising the mean Euclidean reconstruction error can encourage different stochastically

*Full version of a short paper that was published in the Digest of the 5th IMA International Conference on Mathematics in Signal Processing, 18-20 December 2000, Warwick University, UK.

sampled code indices to become associated with different input subspaces [7]. Another advantage is that it is very easy to connect SVQs together, by using the vector of code index probabilities computed by one SVQ as the input vector to another SVQ [8].

In Section 2 various pieces of previously published theory are unified to give a coherent account of SVQs. In Section 3 the results of some new numerical simulations are presented, which demonstrate how the code indices in a SVQ can become associated in various ways with input subspaces. In the appendices various derivations relating to the detailed training of an SVQ are presented.

2 Theory

In this section various pieces of previously published theory are unified to establish a coherent framework for modelling SVQs. In Section 2.1 the basic theory of folded Markov chains (FMC) is given [9], and in Section 2.2 it is extended to the case of high-dimensional input data [10]. Finally, in Section 2.3 the theory is further generalised to chains of linked FMCs [8].

2.1 Folded Markov Chains

The basic building block of the encoder/decoder model used in this paper is the folded Markov chain (FMC) [9]. Thus an input vector x is encoded as a code index vector y, which is then subsequently decoded as a reconstruction x' of the input vector. Both the encoding and decoding operations are allowed to be probabilistic, in the sense that y is a sample drawn from $\Pr(y|x)$, and x' is a sample drawn from $\Pr(x'|y)$, where $\Pr(y|x)$ and $\Pr(x'|y)$ are Bayes' inverses of each other, as given by $\Pr(x'|y) = \frac{\Pr(y|x)\Pr(x)}{\int dz \Pr(y|z)\Pr(z)}$, and $\Pr(x)$ is the prior probability from which x was sampled. Because the chain of dependences in passing from x to y and then to x' is first order Markov (i.e. it is described by the directed graph $(x \longrightarrow y \longrightarrow x')$, and because the two ends of this Markov chain (i.e. x and x') live in the same vector space, it is called a *folded* Markov chain (FMC). The operations that occur in an FMC are summarised in Figure 1.

In order to ensure that the FMC encodes the input vector optimally, a measure of the reconstruction error must be minimised. There are many possible ways to define this measure, but one that is consistent with many previous results, and which also leads to many new results, is the mean Euclidean reconstruction error measure D, which is defined as

$$D \equiv \int dx \Pr(x) \sum_{y_1=1}^{M} \sum_{y_2=1}^{M} \cdots \sum_{y_n=1}^{M} \Pr(y|x) \int dx' \Pr(x'|y) \, ||x - x'||^2 \qquad (1)$$

where $y = (y_1, y_2, \cdots, y_n), 1 \leq y_i \leq M$ is assumed, $\Pr(x)\Pr(y|x)\Pr(x'|y)$ is the joint probability that the FMC has state (x, y, x'), $||x - x'||^2$ is the Euclidean

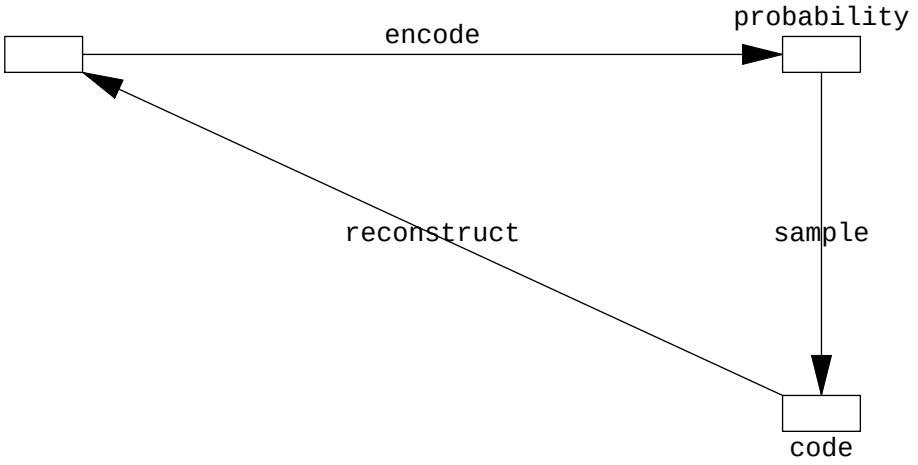

Figure 1: A folded Markov chain (FMC) in which an input vector x is encoded as a code index vector y that is drawn from a conditional probability $\Pr(y|x)$, which is then decoded as a reconstruction vector x' drawn from the Bayes' inverse conditional probability $\Pr(x'|y)$.

reconstruction error, and $\int dx \sum_{y_1=1}^{M} \sum_{y_2=1}^{M} \cdots \sum_{y_n=1}^{M} \int dx'(\cdots)$ sums over all possible states of the FMC (weighted by the joint probability).

The Bayes' inverse probability $\Pr(x'|y)$ may be integrated out of this expression for D to yield

$$D = 2 \int dx \Pr(x) \sum_{y_1=1}^{M} \sum_{y_2=1}^{M} \cdots \sum_{y_n=1}^{M} \Pr(y|x) \, ||x - x'(y)||^2 \tag{2}$$

where the reconstruction vector $x'(y)$ is defined as $x'(y) \equiv \int dx \Pr(x|y)x$. Because of the quadratic form of the objective function, it turns out that $x'(y)$ may be treated as a free parameter whose optimum value (i.e. the solution of $\frac{\partial D}{\partial x'(y)} = 0$) is $\int dx \Pr(x)x$, as required.

It was shown in [9] that the standard VQ [3] and topograpic mappings [11] automatically emerge as special cases when D is minimised. In this approach, topographic mappings emerge as the optimal coding scheme when the code is to be transmitted along a noisy communication channel before being decoded [12, 13].

2.2 High Dimensional Input Spaces

A problem with the standard VQ is that its code book grows exponentially in size as the dimensionality of the input vector is increased, assuming that

the contribution to the reconstruction error from each input dimension is held constant. This means that such VQs are useless for encoding extremely high dimensional input vectors, such as images. The usual solution to this problem is to manually partition the input space into a number of lower dimensional subspaces, and then to encode each of these subspaces separately. However, it would be very useful if this partitioning could be done automatically, in such a way that typically the correlations *within* each subspace were much stronger than the correlations *between* subspaces, so that the subspaces were approximately statistically independent of each other. The purpose of this paper is to present a solution to this problem.

The key step in solving this problem is to constrain the minimisation of D in such a way as to encourage the formation of code schemes in which each component of the code vector y codes a different subspace of the input vector x. There are two related constraints that may be imposed on $\Pr(y|x)$ and $x'(y)$ which may be summarised as

$$\begin{aligned} \Pr(y|x) &= \Pr(y_1|x)\Pr(y_2|x)\cdots\Pr(y_n|x) \\ x'(y) &= \frac{1}{n}\sum_{i=1}^{n} x'(y_i) \end{aligned} \tag{3}$$

Thus each component y_i (for $i = 1, 2, \cdots, n$ and $1 \leq y_i \leq M$) is an *independent* sample drawn from the codebook using $\Pr(y_i|x)$ (which is assumed to be the same function for all i), and the reconstruction vector $x'(y)$ (vector argument) is assumed to be a *superposition* of n contributions $x'(y_i)$ (scalar argument) for $i = 1, 2, \cdots, n$. Taken together, these constraints encourage the formation of coding schemes in which independent subspaces are separately coded, as required.

The constraints in Equation 3 prevent the full space of possible values of $\Pr(y|x)$ or $x'(y)$ from being explored as D is minimised, so they lead to an *upper bound* $D_1 + D_2$ on the FMC objective function D (i.e. $D \leq D_1 + D_2$), which may be derived as [10]

$$\begin{aligned} D_1 &\equiv \tfrac{2}{n}\textstyle\int dx \Pr(x) \sum_{y=1}^{M} \Pr(y|x)\,||x - x'(y)||^2 \\ D_2 &\equiv \tfrac{2(n-1)}{n}\textstyle\int dx \Pr(x) \left|\left| x - \sum_{y=1}^{M} \Pr(y|x)x'(y) \right|\right|^2 \end{aligned} \tag{4}$$

Note that M (size of codebook) and n (number of samples drawn from codebook using $\Pr(y|x)$) are effectively model order parameters, whose values need to be chosen appropriately for each encoder optimisation problem. The properties of the optimum solution depend critically on the interplay between the statistical properties of the training data and the model order parameters M and n, as will be seen in the simulations in Section 3.

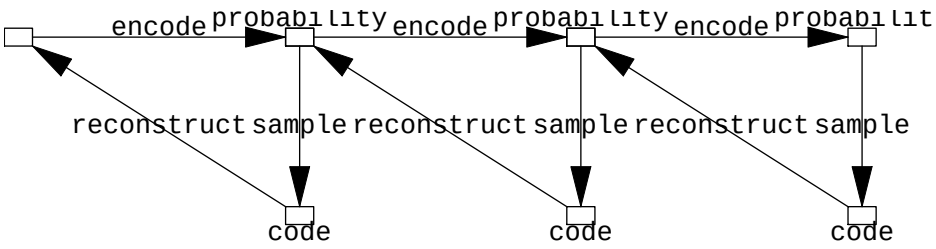

Figure 2: A chain of linked FMCs, in which the output from each stage is its vector of posterior probabilities (for all values of the code index), which is then used as the input to the next stage. Only 3 stages are shown, but any number may be used. More generally, any acyclically linked network of FMCs may be used.

2.3 Chains of Linked FMCs

The FMC illustrated in Figure 1 may be generalised to a chain of linked FMCs as shown in Figure 2. Each stage in this chain is an FMC of the type shown in Figure 1, and the vector of probabilities (for all values of the code index) computed by each stage is used as the input vector to the next stage; there are other ways of linking the stages together, but this is the simplest possibility. The overall objective function is a weighted sum of the FMC objective functions derived from each stage. The total number of free parameters in an L stage chain is $3L-1$, which is the sum of 2 free parameters for each of the L stages, plus $L-1$ weighting coefficients; there are $L-1$ rather than L weighting coefficients because the overall normalisation of the objective function does not affect the optimum solution.

The chain of linked FMCs may be expressed mathematically by first of all introducing an index l to allow different stages of the chain to be distinguished thus

$$\begin{aligned} M &\longrightarrow M^{(l)} \\ x &\longrightarrow x^{(l)} \\ y &\longrightarrow y^{(l)} \\ x' &\longrightarrow x'^{(l)} \\ n &\longrightarrow n^{(l)} \\ D &\longrightarrow D^{(l)} \\ D_1 &\longrightarrow D_1^{(l)} \\ D_2 &\longrightarrow D_2^{(l)} \end{aligned} \tag{5}$$

The stages are then defined and linked together thus

$$\begin{aligned} x^{(l)} &\longrightarrow y^{(l)} \longrightarrow x'^{(l)} \\ x^{(l+1)} &= \left(x_1^{(l+1)}, x_2^{(l+1)}, \cdots, x_{M^{(l)}}^{(l+1)}\right) \\ x_i^{(l+1)} &= \Pr(y^{(l)} = i | x^{(l)}), 1 \leq i \leq M^{(l)} \end{aligned} \tag{6}$$

The objective function and its upper bound are then given by

$$\begin{aligned} D &= \sum_{l=1}^{L} s^{(l)} D^{(l)} \\ &\leq D_1 + D_2 \\ &= \sum_{l=1}^{L} s^{(l)} \left(D_1^{(l)} + D_2^{(l)} \right) \end{aligned} \tag{7}$$

where $s^{(l)} \geq 0$ is the weighting that is applied to the contribution of stage l of the chain to the overall objective function.

3 Simulations

In this section the results of various simulations are presented, which demonstrate some of the types of self-organising behaviour exhibited by an encoder that consists of a chain of linked FMCs. Synthetic, rather than real, training data are used in all of the simulations, because this allows the basic types of behaviour to be cleanly demonstrated.

In Section 3.1 the training data is described. In Section 3.2 a single stage encoder is trained on data that is a superposition of two randomly positioned objects. In Section 3.3 this is generalised to objects with correlated positions, and three different types of behaviour are demonstrated: factorial encoding using both a 1-stage and a 2-stage encoders (Section 3.4), joint encoding using a 1-stage encoder (Section 3.5), and invariant encoding using a 2-stage encoder (Section 3.6).

In Appendix A the derivatives of the objective function are derived, and in Appendix B a gradient descent training algorithm based on these derivatives is presented.

3.1 Training Data

The key property that this type of self-organising encoder exhibits is its ability to automatically split up high-dimensional input spaces into lower-dimensional subspaces, each of which is separately encoded. This self-organisation manifests itself in many different ways, depending on the interplay between the statistical properties of the training data, and the 3 free parameters (i.e. the code book size M, the number of code indices sampled n, and the stage weighting s) per stage of the encoder (see Section 2.3).

In order to demonstrate the various different basic types of self-organisation it is necessary to use synthetic training data with controlled properties. All of the types of self-organisation that will be demonstrated in this paper may be obtained by training a 1-stage or 2-stage encoder on 24-dimensional data (i.e. $M = 24$) that consists of a superposition of a pair of identical objects (with circular wraparound to remove edge effects), such as is shown in Figure 3.

In the simulations presented below, two different methods of selecting the object positions are used: either the positions are statistically independent, or

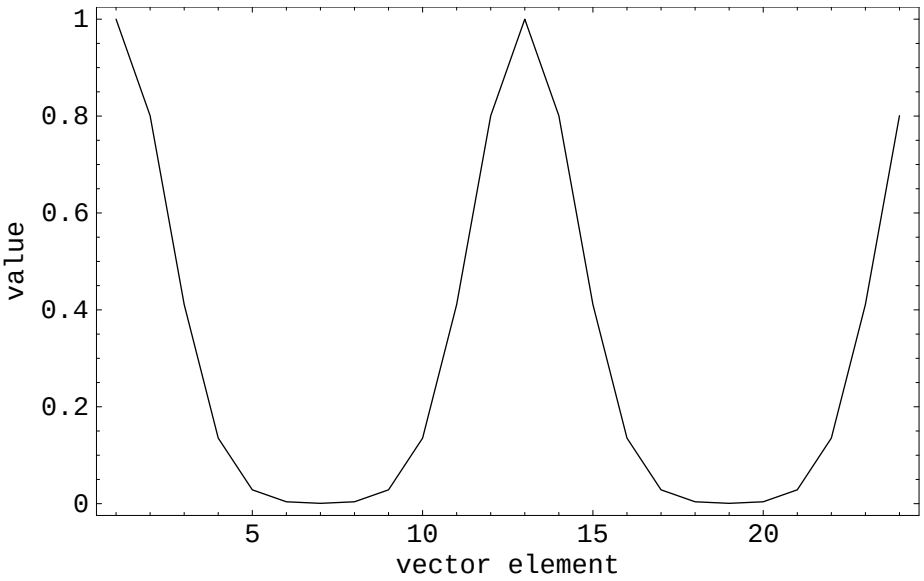

Figure 3: An example of a typical training vector for $M = 24$. Each object is a Gaussian hump with a half-width of 1.5 units, and peak amplitude of 1. The overall input vector is formed as a linear superposition of the 2 objects. Note that the input vector is wrapped around circularly to remove minor edge effects that would otherwise arise.

they are correlated. In the independent case, each object position is a random integer in the interval $[1, 24]$. In the correlated case, the first object position is a random integer in the interval $[1, 24]$, and the second object position is chosen *relative to* the first one as an integer in the range $[4, 8]$, so that the mean object separation is 6 units.

3.2 Independent Objects

The simplest demonstration is to let a single stage encoder discover the fact that the training data consists of a superposition of a pair of objects, which is a type of independent component analysis (ICA) [14]. This may readily be done by setting the parameter values as follows: code book size $M = 16$, number of code indices sampled $n = 20$, $\varepsilon = 0.2$ for 500 training steps, $\varepsilon = 0.1$ for a further 500 training steps.

The self-organisation of each of the 16 reconstruction vectors as training progresses (measured down the page) is shown in Figure 4. After some initial confusion, the reconstruction vectors self-organise so that each code index corresponds to a *single* object at a well defined location, whose width automatically adjusts itself so that the M reconstruction vectors cover the whole input space.

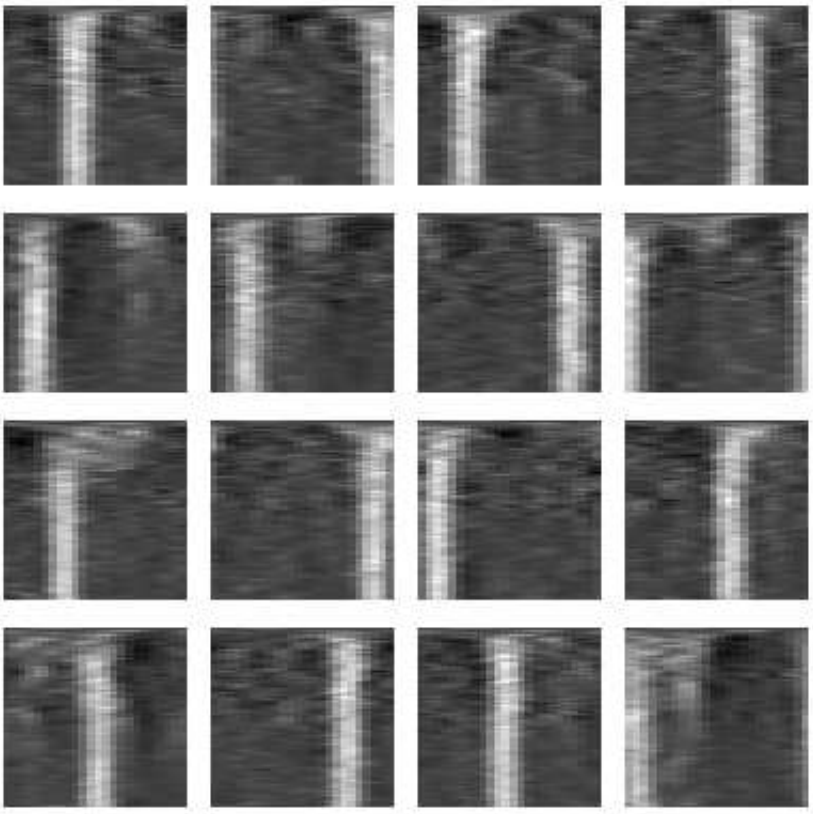

Figure 4: A factorial encoder emerges when a single stage encoder is trained on data that is a superposition of 2 objects in independent locations.

This behaviour is non-trivial, because each training vector is a superposition of a *pair* of objects at independent locations, so two different code index values must be sampled by the encoder (assuming that the two objects are not at the same location); the relatively large choice $n = 20$ ensures that it is highly likely that both code index values will be amongst the n random samples [7]. This result is called a factorial encoder, because the objects are encoded separately.

The case of a joint encoder, where each code index corresponds to a *pair* of objects at well defined locations, requires a rather large code book when the objects are independent. However, when correlations between the objects are introduced then the code book can be reduced to a manageable size, as will be demonstrated in the next section.

3.3 Correlated Objects

If the positions of the pair of objects are mutually correlated, then they can be encoded in 3 fundamentally different ways:

1. Factorial encoder. This encoder ignores the correlations between the objects, and encodes them as if they were 2 independent objects. Each code index thus encodes a single object position, so many code indices must be sampled in order to virtually guarantee that both object positions are encoded [7]. This result is a type of independent component analysis (ICA) [14].

2. Joint encoder. This encoder regards each possible joint placement of the 2 objects as a distinct configuration. Each code index thus encodes a pair of object positions, so only one code index needs to be sampled in order to guarantee that both object positions are encoded [7]. This result is basically the same as what would be obtained by using a standard VQ [3].

3. Invariant encoder. This encoder regards each possible placement of the centroid of the 2 objects as a distinct configuration, but regards all possible object separations (for a given centroid) as being equivalent. Each code index thus encodes only the centroid of the pair of objects. This type of encoder does not arise when the objects are independent. This is similar to self-organising transformation invariant detectors described in [15].

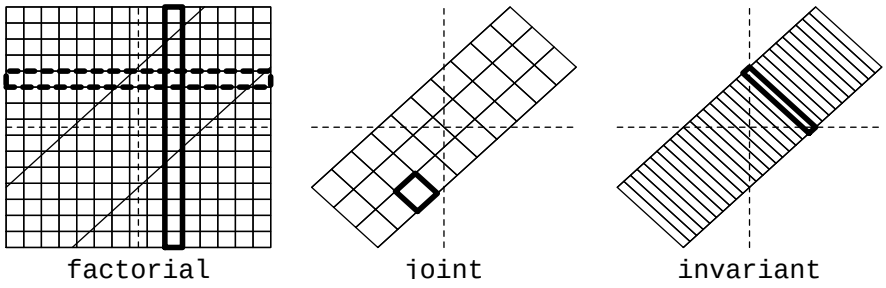

Figure 5: Three alternative ways of using 30 code indices to encode a pair of correlated variables. The typical code cells are shown in bold.

Each of these 3 possibilities is shown in Figure 5, where the diagrams are meant only to be illustrative. The correlated variables live in the large 2-dimensional rectangular region extending from bottom-left to top-right of each diagram.

The factorial encoder has two orthogonal sets of long thin rectangular code cells, and the diagram shows how a pair of such cells intersect to define a small square code cell. The joint encoder behaves as a standard vector quantiser, and is illustrated as having a set of square code cells, although their shapes will not

be as simple as this in practice. The invariant encoder ideally has a set of long thin rectangular code cells that encode only the long diagonal dimension.

In all 3 cases there is overlap between code cells. In the case of the factorial and joint encoders the overlap tends to be only between nearby code cells, whereas in the case of an invariant encoder the range of the overlap is usually much greater, as will be seen in the numerical simulations below. In practice the optimum encoder may not be a clean example of one of the types illustrated in Figure 5, as will also be seen in the numerical simulations below.

3.4 Factorial Encoding

A factorial encoder may be trained by setting the parameter values as follows: code book size $M = 16$, number of code indices sampled $n = 20$, $\varepsilon = 0.2$ for 500 training steps, $\varepsilon = 0.1$ for a further 500 training steps.

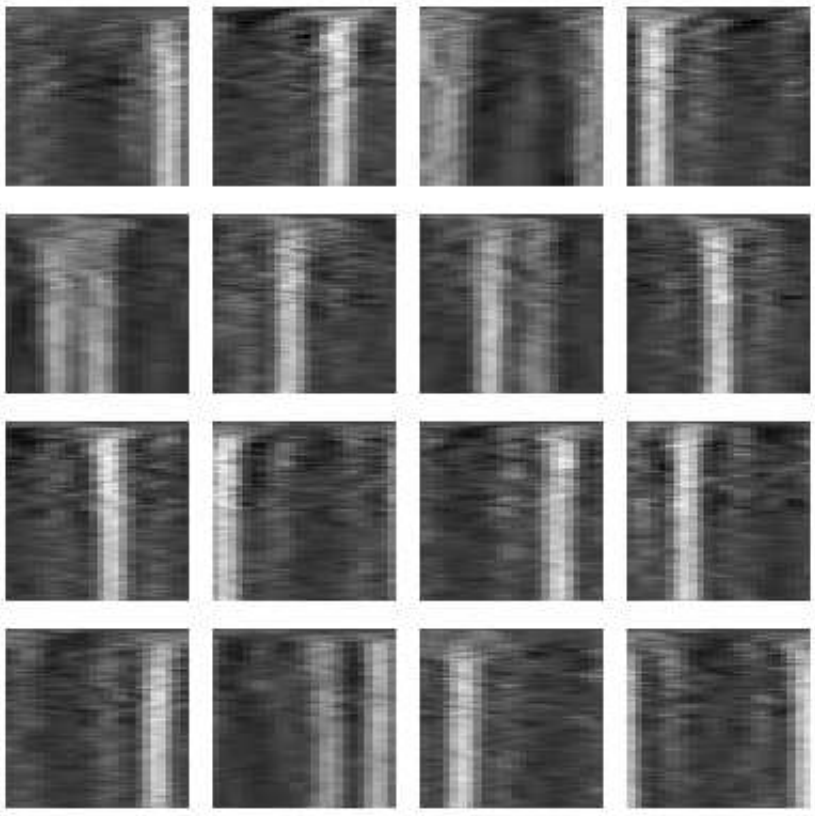

Figure 6: A factorial encoder emerges when a single stage encoder is trained on data that is a superposition of 2 objects in correlated locations.

The result is shown in Figure 6 which should be compared with the result for independent objects in Figure 4. The presence of correlations degrades the quality of this factorial code relative to the case of independent objects. The contamination of the factorial code takes the form of a few code indices which respond jointly to the pair of objects.

The joint coding contamination of the factorial code can be reduced by using a 2-stage encoder, in which the second stage has the same values of M and n as the first stage (although identical parameter values are not necessary), and (in this case) both stages have the same weighting in the objective function (see Equation 7).

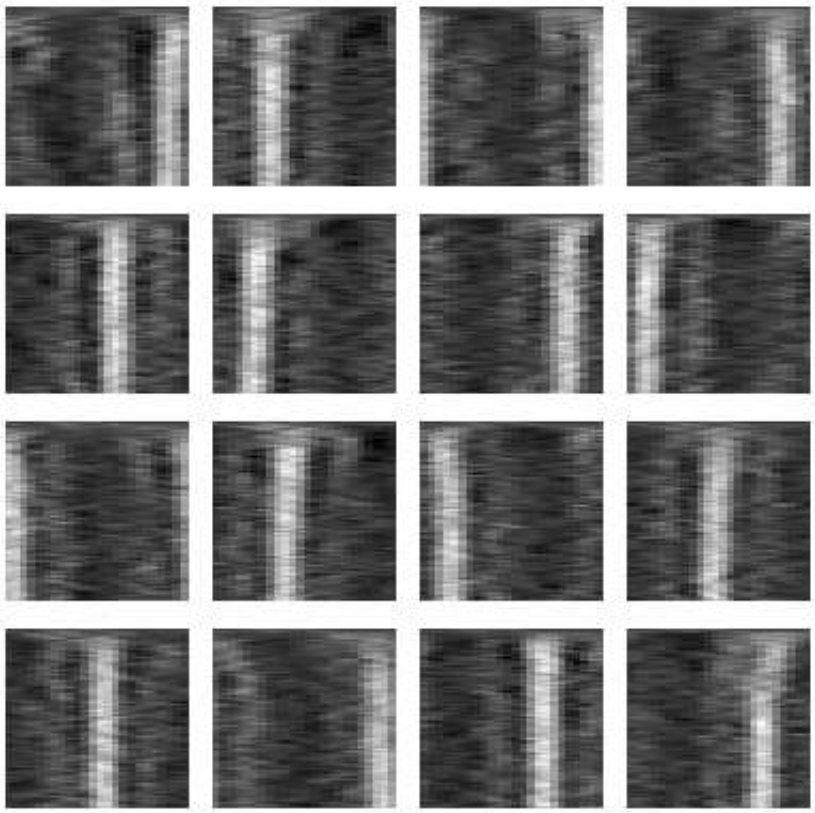

Figure 7: The factorial encoder is improved, by the removal of the joint encoding contamination, when a 2-stage encoder is used.

The results are shown in Figure 7. The reason that the second stage encourages the first to adopt a pure factorial code is quite subtle. The result shown in Figure 7 will lead to the first stage producing an output in which 2 code indices

(one for each object) each typically have probability $\frac{1}{2}$ of being sampled, and all of the remaining code indices have a very small probability (this is an approximation which ignores the fact that the code cells overlap). On the other hand, Figure 6 will lead to an output in which the probability can be concentrated on a single code index, if it can jointly code the pair of objects. However, the contribution of the second stage to the overall objective function encourages it to encode the vector of probabilities output by the first stage with minimum Euclidean reconstruction error, which is easier to do if the situation is as in Figure 7 rather than as in Figure 6. In effect, the second stage likes to see an output from the first stage in which more than one code index has a significant probability of being sampled, which favours factorial coding over joint encoding.

3.5 Joint Encoding

A joint encoder may be trained by setting the parameter values as follows: code book size $M = 16$, number of code indices sampled $n = 3$, $\varepsilon = 0.2$ for 500 training steps, $\varepsilon = 0.1$ for a further 500 training steps, $\varepsilon = 0.05$ for a further 1000 training steps. This is the same as the parameter values for the factorial encoder above, except that n has been reduced to $n = 3$, and the training schedule has been extended.

The result is shown in Figure 8. After some initial confusion, the reconstruction vectors self-organise so that each code index corresponds to a *pair* of objects at well defined locations, so the code index jointly encodes the pair of object positions; this is a joint encoder. The small value of n prevents a factorial encoder from emerging [7].

3.6 Invariant Encoding

An invariant encoder may be trained by using a 2-stage encoder, and setting the parameter values identically in each stage as follows (where the weighting of the second stage relative to the first is denoted as s): code book size $M = 16$, number of code indices sampled $n = 3$, $\varepsilon = 0.2$ and $s = 5$ for 500 training steps, $\varepsilon = 0.1$ and $s = 10$ for a further 500 training steps, $\varepsilon = 0.05$ and $s = 20$ for a further 500 training steps, $\varepsilon = 0.05$ and $s = 40$ for a further 500 training steps. This is basically the same as the parameter values used for the joint encoder above, except that there are now 2 stages, and the weighting of the second stage is progressively increased throughout the training schedule. Note that the large value that is used for s is offset to a certain extent by the fact that the ratio of the normalisation of the inputs to the first and second stages is very large; the anomalous normalisation of the input to the first stage could be removed by insisting that the input to the first stage is a vector of probabilities, but that is not done in these simulations.

The result is shown in Figure 9. During the early part of the training schedule the weighting of the second stage is still relatively small, so it has the effect of turning what would otherwise have been a joint encoder into a factorial encoder; this is analogous to the effect observed when Figure 6 becomes Figure

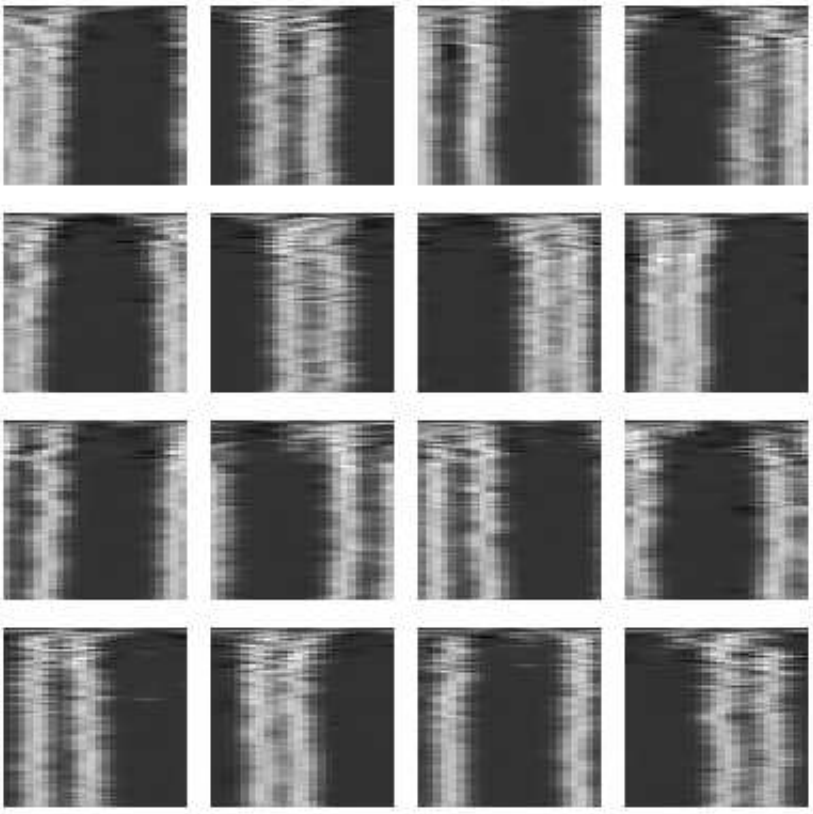

Figure 8: A joint encoder emerges when a single stage encoder is trained on data that is a superposition of 2 objects in correlated locations.

7. However, as the training schedule progresses the weighting of the second stage increases further, and the reconstruction vectors self-organise so that each code index corresponds to a *pair* of objects with a well defined centroid but indeterminate separation. Thus each code index encodes only the centroid of the pair of objects and ignores their separation. This is a new type of encoder that arises when the objects are correlated, and it will be called an *invariant* encoder, in recognition of the fact that its output is invariant with respect to the separation of the objects.

Note that in these results there is a large amount of overlap between the code cells, which should be taken into account when interpreting the illustration in Figure 5. This is an extreme example the second stage preferring an output from the first stage in which more than one code index has a significant probability of being sampled; the large amount of overlap between code cells means that

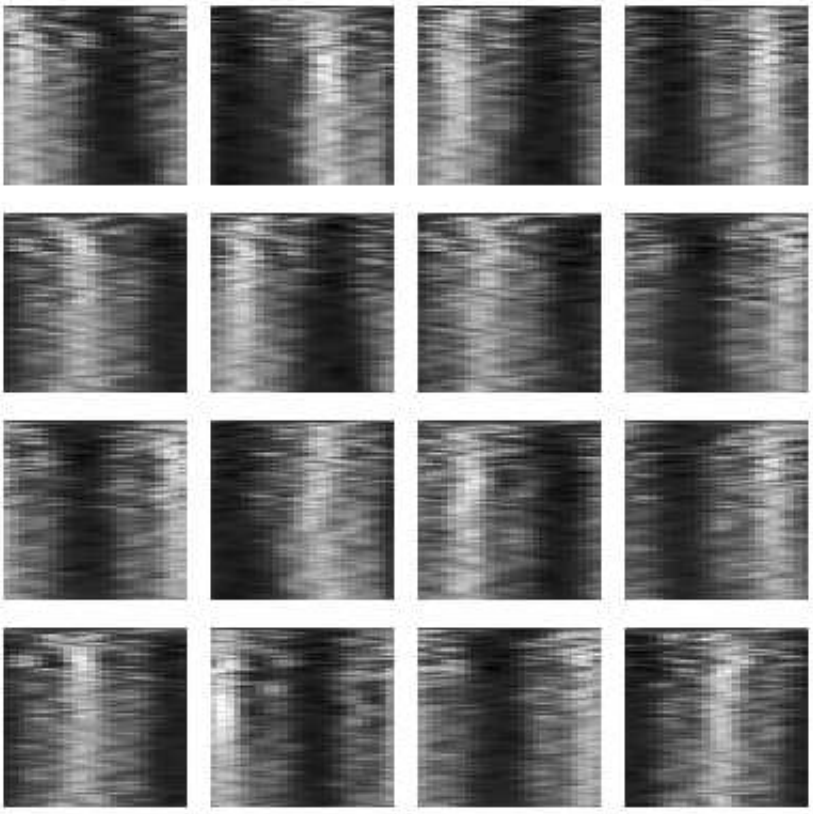

Figure 9: An invariant encoder emerges when 2-stage encoder is trained on data that is a superposition of 2 objects in correlated locations.

many code indices have a significant probability of being sampled.

4 Conclusions

The numerical results presented in this paper show that a stochastic vector quantiser (SVQ) can self-organise to find a variety of different types of way of encoding high-dimensional input vectors. Three fundamentally different types of encoder have been demonstrated, which differ in the way that they build a reconstruction that approximates the input vector:

1. A factorial encoder uses a reconstruction that is superposition of a *number* of vectors that each lives in a well defined input subspace, which is useful for discovering constituent objects in the input vector. This result is a type of independent component analysis (ICA) [14].

2. A joint encoder uses a reconstruction that is a *single* vector that lives in the whole input space. This result is basically the same as what would be obtained by using a standard VQ [3].

3. An invariant encoder uses a reconstruction that is a *single* vector that lives in a subspace of the whole input space, so it ignores some dimensions of the input vector, which is therefore useful for discovering correlated objects whilst rejecting uninteresting fluctuations in their relative coordinates. This is similar to self-organising transformation invariant detectors described in [15].

More generally, the encoder will be a hybrid of these basic types, depending on the interplay between the statistical properties of the input vector and the parameter settings of the SVQ.

A Derivatives of the Objective Function

In order to minimise $D_1 + D_2$ it is necessary to compute its derivatives. The derivatives were presented in detail in [10] for a single stage chain (i.e. a single FMC). The purpose of this appendix is to extend this derivation to a multi-stage chain of linked FMCs. In order to write the various expressions compactly, infinitesimal variations will be used thoughout this appendix, so that $\delta(uv) = \delta uv + u\delta v$ will be written rather than $\frac{\partial(uv)}{\partial\theta} = \frac{\partial u}{\partial\theta}v + u\frac{\partial v}{\partial\theta}$ (for some parameter θ). The calculation will be done in a top-down fashion, differentiating the objective function first, then differentiating anything that the objective function depends on, and so on following the dependencies down until only constants are left (this is essentially the chain rule of differentiation).

The derivative of $D_1 + D_2$ (defined in Equation 7) is given by

$$\delta \sum_{l=1}^{L} s^{(l)} \left(D_1^{(l)} + D_2^{(l)} \right) = \sum_{l=1}^{L} s^{(l)} \left(\delta D_1^{(l)} + \delta D_2^{(l)} \right) \tag{8}$$

The derivatives of the $D_1^{(l)}$ and $D_2^{(l)}$ parts (defined in Equation 4, with appropriate (l) superscripts added) of the contribution of stage l to $D_1 + D_2$ are given by (dropping the (l) superscripts again, for clarity)

$$\begin{aligned}
\delta D_1 &= \tfrac{2}{n} \textstyle\int dx \Pr(x) \sum_{y=1}^{M} \Big(\delta Pr(y|x) \, ||x - x'(y)||^2 \\
&\quad + 2 \Pr(y|x) \left(\delta x - \delta x'(y) \right) \Big) . \left(x - x'(y) \right) \\
\delta D_2 &= \tfrac{4(n-1)}{n} \textstyle\int dx \Pr(x) \left(\delta x - \sum_{y=1}^{M} \left(\delta Pr(y|x) x'(y) \right. \right. \\
&\quad \left. \left. + \Pr(y|x) \delta x'(y) \right) \right) . \left(x - \sum_{y'=1}^{M} \Pr(y'|x) x'(y') \right)
\end{aligned} \tag{9}$$

The first step in modelling $\Pr(y|x)$ is to explicitly state the fact that it is a probability, which is a non-negative normalised quantity. This may be done as

follows

$$\Pr(y|x) = \frac{Q(y|x)}{\sum_{y'=1}^{M} Q(y'|x)} \tag{10}$$

where $Q(y|x) \geq 0$. The $Q(y|x)$ are unnormalised probabilities, and $\sum_{y'=1}^{M} Q(y'|x)$ is the normalisation factor. The derivative of $\Pr(y|x)$ is given by

$$\frac{\delta Pr(y|x)}{\Pr(y|x)} = \frac{1}{Q(y|x)} \left(\delta Q(y|x) - \Pr(y|x) \sum_{y'=1}^{M} \delta Q(y'|x) \right) \tag{11}$$

The second step in modelling $\Pr(y|x)$ is to introduce an explicit parameteric form for $Q(y|x)$. The following sigmoidal function will be used in this paper

$$Q(y|x) = \frac{1}{1 + \exp(-w(y).x - b(y))} \tag{12}$$

where $w(y)$ is a weight vector and $b(y)$ is a bias. The derivative of $Q(y|x)$ is given by

$$\delta Q(y|x) = Q(y|x) \left(1 - Q(y|x)\right) \left(\delta w(y).x + w(y).\delta x + \delta b(y)\right) \tag{13}$$

This has reduced the δD_1 and δD_2 derivatives to $\delta w(y)$, $\delta b(y)$, $\delta x'(y)$ and δx derivatives. The $\delta w(y)$, $\delta b(y)$ and $\delta x'(y)$ derivatives relate directly to the parameters being optimised and thus need no further simplification, however the δx derivatives in Equation 9 and Equation 13 need some further attention. The δx derivative arises only in multi-stage chains of FMCs, and because of the way in which stages of the chain are linked together (see Equation 6) it is equal to the derivative of the vector of probabilities output by the previous stage. Thus the δx derivative may be obtained by following its dependencies back through the stages of the chain until the first layer is reached; this is essentially the chain rule of differentiation. This ensures that for each stage the partial derivatives include the additional contributions that arise from forward propagation through later stages, as described in Appendix B.

B Training Algorithm

Assuming that $\Pr(y|x)$ is modelled as in appendix A (i.e. $\Pr(y|x) = \frac{Q(y|x)}{\sum_{y'=1}^{M} Q(y'|x)}$ and $Q(y|x) = \frac{1}{1+\exp(-w(y).x-b(y))}$), then the partial derivatives of $D_1 + D_2$ with respect to the 3 types of parameters in a single stage of the encoder may be denoted as

$$\begin{aligned} g_w(y) &\equiv \tfrac{\partial(D_1+D_2)}{\partial w(y)} \\ g_b(y) &\equiv \tfrac{\partial(D_1+D_2)}{\partial b(y)} \\ g_x(y) &\equiv \tfrac{\partial(D_1+D_2)}{\partial x'(y)} \end{aligned} \tag{14}$$

This may be generalised to each stage of a multi-stage encoder by including an (l) superscript, and ensuring that for each stage the partial derivatives include the additional contributions that arise from forward propagation through later stages; this is essentially an application of the chain rule of differentiation, using the derivatives $\frac{\partial x^{(l+1)}}{\partial w^{(l)}(y^{(l)})}$ and $\frac{\partial x^{(l+1)}}{\partial b^{(l)}(y^{(l)})}$ to link the stages together (see appendix A).

A simple algorithm for updating these parameters is (omitting the (l) superscript, for clarity)

$$\begin{aligned} w(y) &\longrightarrow w(y) - \varepsilon \frac{g_w(y)}{g_{w,0}} \\ b(y) &\longrightarrow b(y) - \varepsilon \frac{g_b(y)}{g_{b,0}} \\ x'(y) &\longrightarrow x'(y) - \varepsilon \frac{g_x(y)}{g_{x,0}} \end{aligned} \tag{15}$$

where ε is a small update step size parameter, and the three normalisation factors are defined as

$$\begin{aligned} g_{w,0} &\equiv \max_y \sqrt{\frac{||g_w(y)||^2}{\dim x}} \\ g_{b,0} &\equiv \max_y |b(y)| \\ g_{x,0} &\equiv \max_y \sqrt{\frac{||g_x(y)||^2}{\dim x}} \end{aligned} \tag{16}$$

The $\frac{g_w(y)}{g_{w,0}}$ and $\frac{g_x(y)}{g_{x,0}}$ factors ensure that the maximum update step size for $w(y)$ and $x'(y)$ is $\varepsilon \dim x$ (i.e. ε per dimension), and the $\frac{g_b(y)}{g_{b,0}}$ factor ensures that the maximum update step size for $b(y)$ is ε. When a stationary point of $D_1 + D_2$ is reached, the finite size of ε prevents the parmater values from converging to a perfectly stationary solution, and instead they jump around in its neighbourhood.

This update algorithm can be generalised to use a different ε for each stage of the encoder, and also to allow a different ε to be used for each of the 3 types of parameter. Furthermore, the size of ε can be varied as training proceeds, usually starting with a large value, and then gradually reducing its size to obtain an accurate estimate of the stationary solution. It is not possible to give general rules for exactly how to do this, because training conditions depend very much on the statistical properties of the training set.

References

[1] Gray, R. M. (April 1984). *Vector quantisation. IEEE Acoust., Speech, Signal Processing Mag.*, 4–29.

[2] Gersho, A., & Gray, R. M. (1992). *Vector quantisation and signal processing.* Kluwer.

[3] Linde, Y., Buzo, A., & Gray, R. M. (1980). *An algorithm for vector quantiser design. IEEE Trans. COM*, **28**(1), 84–95.

[4] Zeger, K., & Gersho, A. (1989). *Stochastic relaxation algorithm for improved vector quantiser design. Electronics Lett.*, **25**(14), 896–898.

[5] Yair, E., Zeger, K., & Gersho, A. (1992). *Competitive learning and soft competition for vector quantiser design. IEEE Transactions on Signal Processing*, **40**(2), 294–309.

[6] Torres, L., Casas, J. R., & Arias, E. (1997). *Stochastic vector quantisation of images. Signal Process.*, **62**(3), 291–301.

[7] Luttrell, S. P. (1999). *Self-organised modular neural networks for encoding data.* In A. J. Sharkey (Ed.), *Combining Artificial Neural Nets: Ensemble and Modular Multi-Net Systems* (pp. 235–263). Springer-Verlag.

[8] Luttrell, S. P. (1999). *An adaptive network for encoding data using piecewise linear functions.* In *Proceedings of the 9th International Conference on Artificial Neural Networks (ICANN99)* (pp. 198–203). Edinburgh.

[9] Luttrell, S. P. (1994). *A Bayesian analysis of self-organising maps. Neural Computation*, **6**(5), 767–794.

[10] Luttrell, S. P. (1997). *A theory of self-organising neural networks.* In S. W. Ellacott, J. C. Mason & I. J. Anderson (Ed.), *Mathematics of Neural Networks: Models, Algorithms and Applications* (pp. 240–244). Kluwer.

[11] Kohonen, T. (1984). *Self-organisation and associative memory.* Springer-Verlag.

[12] Farvardin, N. (1990). *A study of vector quantisation for noisy channels. IEEE Transactions on Information Theory*, **36**(4), 799–809.

[13] Kumazawa, H., Kasahara, M., & Namekawa, T. (1984). *A construction of vector quantisers for noisy channels. Electronic Engineering Japan*, **67B**(4), 39–47.

[14] Hyvärinen, A. (1999). *Survey on independent component analysis. Neural Computing Surveys*, **2**, 94–128.

[15] Webber, C. J. (1994). *Self-organisation of transformation-invariant detectors for constituents of perceptual patterns. Network: Computation in Neural Systems*, **5**(4), 471–496.